arXiv:1607.06344v2 [cs.CG] 17 Apr 2017

Solving Equations and Optimization Problems with Uncertainty*

Peter Franek, Marek Krčál, Hubert Wagner

IST Austria

Abstract

We study the problem of detecting zeros of continuous functions that are known only up to an error bound, extending the theoretical work of [26] with explicit algorithms and experiments with an implementation.[1] Further, we show how to use the algorithm for approximating worst-case optima in optimization problems in which the feasible domain is defined by the zero set of a function $f : X \to \mathbb{R}^n$ which is only known approximately.

The algorithm first identifies a subdomain A where the function f is provably non-zero, a simplicial approximation $f' : A \to S^{n-1}$ of $f/|f|$, and then verifies non-extendability of f' to X to certify a zero. Deciding extendability is based on computing the cohomological *obstructions* and their persistence. We describe an explicit algorithm for the *primary and secondary obstruction*, two stages of a sequence of algorithms with increasing complexity. Using elements and techniques of persistent homology, we quantify the persitence of these obstructions and hence of the robustness of zero.

We provide experimental evidence that for random Gaussian fields, the *primary obstruction*—a much less computationally demanding test than the secondary obstruction—is typically sufficient for approximating robustness of zero.

1 Introduction

Motivation. Detecting zeros of $\mathbb{R}^n$-valued functions is equivalent to solving systems of real equations, a fundamental problem of mathematics and

*The research leading to these results has received funding from Austrian Science Fund (FWF): M 1980 and from the Czech Science Foundation (GACR) grant number 15-14484S with institutional support RVO:67985807.

[1] `https://bitbucket.org/robsatteam/rob-sat`

theoretical computer science. Our research is motivated by practical applications, in which the data is often known only approximately. We address the case where the input data is limited to the *approximate* values of a continuous function f with values in $\mathbb{R}^n$, sampled over a finite point set. This uncertainty is handled in a deterministic way: we aim at verifying that *each* continuous function compatible with our partial knowledge of f, has a zero.

Functions that are known only approximately appear in various contexts and are handled in different ways. For example, rounding errors in floating-point computations are systematically treated by methods of *interval arithmetic* and detection of zeros resistant to bounded errors is a frequent problem in this field [31, 22, 3, 34, 7]. Other instances of uncertain functions come from measurements of physical quantities, such as in medical imaging [12, 32, 21] or robotics [41, 6]. We suppose that potential applications include robust detection of level sets $f^{-1}(a)$ in medical image processing, analysing robot trajectory based on data obtained from sensors [6, 7], or computing the *inner approximation* of reachable regions of a robotic arm [33]. The algorithm could also be exploited for analysis of functions obtained by regression (say, in machine learning), where the function is chosen to fit some given set of sampled values.

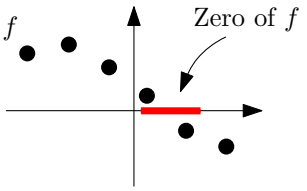

Figure 1: For scalar valued function, existence of a zero can be verified via the intermediate value theorem.

To verify that a function f, of which we only have a limited knowledge, has a zero, is equivalent to showing that *each* potential candidate g for f has a zero. If we only have access to sampled values of f and a Lipschitz constant, then the set of all such admissible functions g is huge and can not be finitely parametrized. However, methods of computational homotopy theory can be applied: the closely related problem of verifying that each continuous r-perturbation of a given function has a zero, can be reduced to the topological extension problem for maps into a sphere [26]. The latter problem can be addressed via means of obstruction theory, using an algorithmic construction of Postnikov towers. Such construction have never been implemented and is in its full generality probably out of reach, given the limitations of computer power. Using a number of simplifications as well as some methods of persistent homology, we present a partial solution to the above problem accompanied by an implementation, complexity analysis and several computational experiments.

Statement of the results. We present an algorithm for detecting zeros

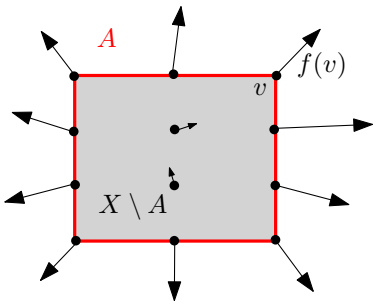

Figure 2: This illustrate a function $f : X \to \mathbb{R}^2$ such that $X = [-1, 1]^2$, $A = \partial X$, and $f : A \to \mathbb{R}^2 \setminus \{0\}$ is homotopic to the identity map. Then the non-extendability of $f|_A$ to $X \to \mathbb{R}^2 \setminus \{0\}$ implies the existence of a zero in $X \setminus A$. We can bound the robustness of zero in $X \setminus A$ by $\min_{v \in A} |f(v)|$ from below.

of vector valued functions $f : X \to \mathbb{R}^n$ on a compact space X and for approximating the *robustness* of zero, that is, a maximal real number $r > 0$ such that every continuous $g : X \to \mathbb{R}^n$ satisfying $\|g - f\| \leq r$ has a zero. By $\|f\|$ we denote the max norm $\max_{x \in X} |f(x)|$ where $|\cdot|$ is a fixed ℓ_p norm in $\mathbb{R}^n$. Nontrivial cases happen if $\dim X \geq n$, as otherwise arbitrarily small perturbations of f avoid zero. For computer representation we assume that the space X is a simplicial complex. The map $f \colon X \to \mathbb{R}^n$ is specified by its values on the vertices which are assumed to be rational, and by a rational value $\alpha > 0$ such that $|f(x) - f(y)| \leq \alpha$ for arbitrary points x and y of any simplex of X. We emphasize that the precise knowledge of f is not needed. The algorithm computes a number $r_1 \in \mathbb{R}$ such that

- Every continuous g, $\|g - f\| \leq r_1$, has a zero.

A positive $r_1 > 0$ is then a certificate of existence of zero of f: we will say that f has an r_1-*robust zero*. Otherwise the algorithm outputs a negative number and gives no guarantee of the existence of zero. Under the dimensional constraints $\dim X \leq n+1$ or $n < 3$, it also computes a number $r_2 > r_1$ such that

- Some continuous g, $\|g - f\| \leq r_2$, has no zero.

Under this dimensional constraint, the gap $r_2 - r_1$ provably converges to zero, if the constant α (and hence our lack of knowledge of f) goes to zero.

The main ingredient is to find a subdomain $A \subseteq X$ where f is provable nonzero and where our knowledge of f is sufficient to determine the homotopy class of $f|_A$ as a map to $\mathbb{R}^n \setminus \{0\} \simeq S^{n-1}$. Then non-extendability to $X \to S^{n-1}$ is a certificate of zero. The primary obstruction measures non-extendability to the n-skeleton of X and the secondary obstruction the non-extendability to the $(n+1)$-skeleton. The constraint $\dim X \leq n+1$ could be generalized, if we implemented oracles for computing higher obstructions, such as discussed in [26, 19].

Our second result is based on computational experiments with random functions. A natural informal question is

- *How typical are functions for which higher obstructions are needed for detecting a zero?*

An example of a function with nontrivial secondary obstruction is any map f from a 4-ball B^4 to $\mathbb{R}^3$ such that $f|_{\partial B^4}$ is homotopic to the Hopf map $S^3 \to S^2$ and hence cannot be extended to $B^4 \to \mathbb{R}^3 \setminus \{0\}$. Such property can be verified, if we are given a sample of function values and a Lipschitz constant. Moreover, the homotopy class of $f|_{\partial B^4}$ does not change, if we slightly perturb these sampled function values.

Surprisingly, when performing experiments with random functions (mainly random Gaussian fields), we observed that higher obstructions are typically *not* needed. Whenever we detected a zero of a randomly generated uncertain function, it was via means of primary obstruction only. We performed experiments with various random functions from a triangulated 4-cube or from a 4-torus into $\mathbb{R}^3$ as well as from a 5-torus into $\mathbb{R}^4$. This observation—if confirmed by theory or by more experiments in different settings—could justify the usage of *only the primary obstruction* in potential future engineering applications.

State of the art. Algorithms for detecting zeros used in software packages are based on iterative methods which are often applicable if f is given by formulas and is differentiable. However, these algorithms usually give no guarantees of correctness: the satisfiability of $f(x) = 0$ is undecidable for any class of real functions f that contain polynomials and the sine function [49].

A number of methods has been proposed for testing the (non-)existence of zeros of continuous functions, exploiting tools ranging from iterative methods in numerical analysis to topology. The problem has been most studied in the case $\dim X = n$. If B^n is a unit ball in $\mathbb{R}^n$, then verifying zeros of $f : B^n \to \mathbb{R}^n$ is equivalent to verifying a fixed point of $f + \mathrm{id}$: here the Brouwer fixed point theorem can be applied [43]. Other methods for zero

verification were studied in the field of interval arithmetic, such as *Miranda's test* [4], *Borsuk's test* [31] and the *degree test* [28]. All of these tests have topological flavour and are stable with respect to perturbations of the input function. Authors of [29] show that the *degree test* can detect a zero of $f : B^n \to \mathbb{R}^n$ whenever the zero of f is *robust* (that is, each g close enough to f has a zero). The above mentioned primary obstruction directly reduces to the degree test if $\dim X = n$. While the topological degree computation has been explicitly described in the literature and also implemented [27], the problem is far more complicated if the domain X has larger dimension than n. In [26] we showed that the existence of a robust zero of piecewise linear functions $X \to \mathbb{R}^n$ is undecidable if $n \geq 3$ is a fixed odd integer and the input X is an arbitrary $(2n-2)$-dimensional simplicial complex.

Zero sets of functions with inherent uncertainty were studied via means of computational topology in the context of *well groups* [11]. In their general settings, well groups associated to $f : X \to Y$ and a subspace $Y' \subseteq Y$ describe properties of the preimage $f^{-1}(Y')$ which persist if we perturbate the input function f. In the important case of $Y = \mathbb{R}^n$ and $Y' = \{0\}$, well groups describe zero sets of functions: namely, the zeroth well groups measures robustness of *existence* of zero and higher well groups reflect further topological properties of the zero sets. In [25] we have shown that the primary obstruction can be used to compute certain subgroup of the well group that in many cases coincides with the full well group (see [25, Thm. 1.4]). A general algorithm for well group computation cannot be expected because the above mentioned undecidability result [26] directly transfers to well groups when $\dim X \geq 2n-2$ as well. Our implementation can be thought of as an approximation of the zeroth well group that extends the work of [20] where the special case $\dim X = n$ is solved.

An application of our algorithm is in worst-case analysis of optimization problems where the feasible domain is defined by equations. Worst-case approach in robust optimization has been widely studied, see [10, 9, 13, 14]. Usually, the uncertainty applies to a finite number of parameters which are assumed to be taken from a known domain. In our approach, we rather work with the space of *all continuous functions* that are compatible with our partial knowledge of f.

Outline and organization of the paper. In the algorithm, we first create a filtration $\{A_r \supseteq A_s\}_{r \leq s}$ of subcomplexes that "approximate" the topological spaces $A(r) := \{x \in X : |f(x)| \geq r\}$. We compute a simplicial approximation $f' : A_r \to \Sigma$ of f where Σ is a given triangulation of the $(n-1)$-sphere. Then we ask for the smallest r such that the restriction of

f' to A_r can be extended to all of X, and show that the robustness of zero of the original function is α-far from r.

Such extendability is decidable if $\dim X \leq 2n-3$ [18], but the only procedere for this we are aware of is based on the algorithm for computing stages of Postnikov towers from [19] that depends on several other papers [38, 17, 48] and is unlikely to be fully implemented in near future. Instead of that, we implemented a persistent version of both the *primary* and *secondary obstruction*, which test extendability to the n- and $(n+1)$-skeleton of X.[2] First, we compute the minimal r_1 for which the cohomological obstructions to extending $f'|_{A_{r_1}}$ to $A_{r_1} \cup X^{(n)}$ (*primary obstruction*) vanishes. Similarly, we compute a minimal $r_2 \geq r_1$ for which $f'|_{A_{r_2}}$ is extendable to $A_{r_2} \cup X^{(n+1)}$ (vanishing of the *secondary obstruction*): this requires to parametrize all extensions to the n-skeleton.

In Section 2 we show how to approximate the spaces $A(r)$ by simplicial complexes and sphere-valued maps $f/|f|$ via simplicial maps. A high-level description of our algorithm is in Sections 3 and 4, with a partial lower-level description in Appendix A and C. In Section 5 we show how to use the method for approximating worst-case optima in optimization problems where the feasible domain is defined by equations. In Section 6 we present some computational experiments with random Gaussian fields. More details about testing and performance are delegated to Appendix D. The last section contains a theoretical worst-case complexity bounds.

2 Discretizing the function f.

In this section we show how to convert the "unknown" continuous function f to its discrete simplicial approximation.

Definition 2.1. *A* continuous filtration *of spaces is a family $(A_r)_{r\in\mathbb{R}}$ such that $A_r \supseteq A_s$ whenever $r \leq s$. A continuous filtration $(A_r)_{r\in\mathbb{R}}$ is called* step-like *whenever there exists a sequence of numbers $-\infty =: r_{-1} < r_0 \leq r_1 \leq r_2 \leq \ldots \leq r_k$ such that for any $r, s \in (r_i, r_{i+1}]$, $A_r = A_s$ holds for all i. Continuous filtrations $(A_r)_r$ and $(B_r)_r$ are called* α-interleaved *whenever $B_{r+\alpha} \subseteq A_r$ and $A_{r+\alpha} \subseteq B_r$ for each $r \in \mathbb{R}$.*

Definition 2.2. *Let $f\colon X \to \mathbb{R}^n$ be a continuous map on a simplicial complex X and let $|\cdot|$ be a norm on $\mathbb{R}^n$.*

[2]The only exception is the case $n = 3$, $\dim X > 3$ where the triviality of secondary obstruction is undecidable in general. However, if X is assumed to be a triangulation of the cube $[0,1]^4$, then our algorithm works with no essential changes. For many other fixed 4-dimensional spaces X the problem is decidable too.

1. *By A_r we denote the* subcomplex *of X spanned by the vertices v of X with $|f(v)| \geq r$.*

2. *By $A(r)$ we denote the* subspace *of X defined by $A(r) = \{x \in X \colon |f(x)| \geq r\}$.*

3. *We say that f is* simplexwise α-Lipschitz *whenever $|f(x) - f(y)| \leq \alpha$ for each pair of points $x, y \in \Delta$ of any simplex $\Delta \in X$.*

Spaces A_r form a step-like filtration where the steps occur for each r equal to $|f(v)|$ for some vertex v of X. Let Σ^{n-1} be the simplicial model of the $(n-1)$-sphere obtained from the boundary of a cross-polytope. A natural sphere-valued approximation of f is then given by the map f' as follows.

Definition 2.3. *Let $f : X \to \mathbb{R}^n$ and V be a subset of the vertices of X. We define the* vertex approximation $f' : V \to \{e_1, -e_1, \ldots, e_n, -e_n\}$ *to by the map that to a vertex v assigns $s_j e_j$, where j is the index of the component of $f(v)$ with* largest *absolute value and s_j is the sign of $f_j(v)$.*[3]

Lemma 2.4. *Let $f \colon X \to \mathbb{R}^n$ be a simplexwise α-Lipschitz map for some constant $\alpha > 0$, $p \in [1, \infty]$ and $|\cdot|$ be the ℓ_p-norm. Then the following holds:*

1. *The continuous filtrations $(A_r)_{r \in \mathbb{R}}$ and $A(r)_{r \in \mathbb{R}}$ are α-interleaved.*

2. *If $r > \alpha n^{1/p}/2$, the vertex approximation $f' : V(A_r) \to V(\Sigma^{n-1})$ defines a simplicial map $f' \colon A_r \to \Sigma^{n-1}$ (that is, it maps simplices to simplices).*

3. *If $r > \alpha n^{1/p}$, then $f' \colon A_r \to \Sigma^{n-1} \subseteq \mathbb{R}^n \setminus \{0\}$ is homotopic to $f|_{A_r} \colon A_r \to \mathbb{R}^n \setminus \{0\}$.*

The simplicial map $f' \colon A_r \to \Sigma^{n-1}$ as above will be called the *simplicial approximation of $f|_{A_r}$*. We note that the choice of the sphere model and the discretization f' of f is independent of the rest of the algorithm given in the following chapters and is not the only possible choice: however, we couldn't find one with approximative properties better than in Lemma 2.4.

Proof of part 1. If a point $x \in |X|$ is not in A_r, we know that $|f(v)| < r$ for some vertex v of a simplex Δ supporting x. The simplexwise α-Lipschitz

[3]For example, if $f(v) = [2, -3]$, then we choose $f'(v) = -e_2$. If there are more components of $f(v)$ with the same absolute value, we choose one by arbitrary rule.

property then implies that $|f(x)| < r + \alpha$, hence $x \notin A(r + \alpha)$. This proves $A(r + \alpha) \subseteq A_r$.

The other inclusion holds because once a point x of a simplex $\Delta \in X$ is in a given A_r, then for arbitrary vertex v of Δ holds $|f(v)| \geq r$ and $|f(x) - f(v)| \leq \alpha$. Thus $|f(x)| \geq r - \alpha$, hence $x \in A(r - \alpha)$. □

Proof of part 2. We want to prove that no adjacent vertices u and v are mapped by f' to $e_i, -e_i$ for some i. Without loss of generality we can assume that $f'(v) = e_1$. Assuming $|f(v)| \geq r$ and the definition of f', the first component (with largest absolute value) must satisfy $f_1(v) \geq rn^{-1/p}$. Similarly if $f'(u) = -e_1$, then $f(u)_1 \leq -rn^{-1/p}$, but this would imply that $\alpha \geq |f(u) - f(v)| \geq 2rn^{-1/p}$, contradicting the assumption $r > \alpha n^{1/p}/2$. □

Proof of part 3. We show that the simplexwise straight-line homotopy between f and f' has values in $\mathbb{R}^n \setminus \{0\}$. Let $\Delta \in X$, v be a vertex of Δ, $x \in \Delta$, and assume that $r > \alpha n^{1/p}$. Again, assume WLOG that $f'(v) = e_1$. Then $f_1(x) \geq f_1(v) - \alpha \geq rn^{-1/p} - \alpha > 0$ and the straightline homotopy between maps $f, f' \colon A_r \to \mathbb{R}^n \setminus \{0\}$ has the first coordinate positive, hence it avoids zero.

□

3 The algorithm using an oracle for persistence of obstructions.

In this section we describe a high-level description of our algorithm for approximating robustness of zero. The specification is as follows:

Input:

- X, a simplicial complex
- $f : X^{(0)} \to \mathbb{R}^n$, function values in vertices
- $\alpha > 0$[4]

Output:

- a lower bound on the robustness of zero (possibly negative)
- an upper bound on the robustness of zero (possibly ∞)

[4]For computability purposes, we assume that $f(v)$ and α are rational or computable.

The unknown function f is thus represented by function values in vertices and a simplex-wise Lipschitz constant α. A negative lower bound or infinite upper bound give no information at all: however, in case of $\dim X \leq n+1$ or $n < 3$, the lower and upper-bounds on robustness will be at most 2α-far from each other. Even outside this dimension range, the computed lower bound on robustness will be at most α-far from the robustness of zero of $f|_{A \cup X^{(n+1)}}$.

Definition 3.1. *Let $X \supseteq A_{r_0} \supseteq A_{r_1} \supseteq \ldots$ be a filtration of simplicial complexes, $r_i \leq r_j$ for $i \leq j$ and $f' : A_{r_0} \to S^{n-1}$. Then the* persistence of primary obstruction *is the smallest r_j such that the restriction of f' to A_{r_j}* can *be extended to a (not necessarily simplicial) map $A_{r_j} \cup X^{(n)} \to S^{n-1}$ where $X^{(n)}$ is the n-skeleton of X. The* persistence of secondary obstruction *is the smallest r_k such that the restriction of f' to A_{r_k} can be extended to $A_{r_k} \cup X^{(n+1)} \to S^{n-1}$.*

In what follows, assume that an oracle is given that, for a filtration of simplicial complexes and a simplicial map f', computes the persistence of secondary obstrucion. We assume that we are given a continuous map $f\colon X \to \mathbb{R}^n$ by its values on the vertices of X, its simplexwise Lipschitz constant α and a norm ℓ_p on $\mathbb{R}^n$ for $p \in [1, \infty]$. The outline of the algorithm follows.

A. (a) Label the set

$$\{|f(v)|\colon v \in V(X) \text{ such that } |f(v)| \geq \alpha n^{1/p}\}$$

by $\{r_0, r_1, \ldots, r_h\}$ so that $r_i \leq r_j$ for $i \leq j$.

(b) For any simplex $\Delta \in X$ compute its filtration value $r(\Delta)$ by

$$r(\Delta) := \min_{v \text{ vertex of } \Delta} |f(v)|. \tag{1}$$

This yields a filtration $A_{r_0} \supseteq \ldots \supseteq A_{r_h}$ that together with the values $r_0, \ldots, r_h$ determines the step-like continuous filtration $(A_r)_{r \in \mathbb{R}}$ from Definition 2.1.

(c) For vertices v of X with $|f(v)| \geq r_0$ compute the vertex approximation $f'(v)$ via Definition 2.3.

B. Use the oracle to compute the persistence of secondary obstruction r_k (Def. 3.1).

(a) **If** $k > 0$: output "robustness of zero is at least $r_k - \alpha$."
Else: output "no guarantee of zero"

(b) **If** $\dim X \leq n+1$ or $n < 3$: output "robustness of zero is at most $r_k + \alpha$"
Else: output "no guarantee of upper bound (robustness of zero is at most ∞)"

The constraints in B(b) could be replaced by $\dim X \leq n-1+k$, if we used an oracle for persistence of the first k obstructions. However, implementing such oracle is theoretically possible only if $k < n-1$. In fact, even the special case $\dim X = 4$ and $n = 3$ is beyond this bound and we cannot implement the oracle for secondary obstruction for this dimension pair with no restrictions on X. In the important special case when X is topologically a cube $[0,1]^4$ and $n = 3$, the general algorithm works with no essential changes. More details about the implementation of the oracle for this dimension pair are given in Appendix A, p. 28.

Theorem 3.2. *The above algorithm outputs correct statements.*

Proof. In [26, Lemma 3.3] we showed that f has an r-robust zero iff $f|_{A(r)}$ is not extendable to a nowhere zero function on X.

Correctness of B(a) Assume first that $k > 0$. That means that $f'|_{A_r}$ is not extendable to a function $X \to \mathbb{R} \setminus \{0\}$ for $r < r_k$. For any $r \in (r_0, r_k)$, $f'|_{A_r}$ is well defined and homotopic to $f|_{A_r} : A_r \to \mathbb{R} \setminus \{0\}$ by Lemma 2.4, 3. It follows that $f|_{A_r}$ is not extendable to a nowhere zero function $X \to \mathbb{R}^n \setminus \{0\}$ either. Further, by Lemma 2.4, 1, $A_r \subseteq A(r-\alpha)$ and the non-extendability of $f|_{A_r}$ immediately implies the non-extendability of $f|_{A(r-\alpha)}$. Thus we conclude that f has at least $(r-\alpha)$-robust zero for arbitrary $r \in (r_0, r_k)$. The set $\{r > 0 \mid \text{each } g,\ \|g - f\| \leq r,\ \text{has a zero}\}$ is closed (see [26, p. 9] for an argument) and hence f has at least $(r_k - \alpha)$-robust zero.

Correctness of B(b) The restriction of f' to A_{r_k} is extendable to $A \cup X^{(n+1)}$, so for any $r \geq r_k$, $f'|_{A_r}$ is extendable to $A_r \cup X^{(n+1)}$ as well. If $\dim X \leq n+1$, then this is equivalent to the extendability to all of X. The cases $n < 3$ reflect low dimensional phenomena: we will show that then the extendability to $A \cup X^{(n)}$ implies the extendability to all of X. If $n = 1$, f' has values in the 0-sphere $S^0 \in \{+, -\}$ and if it can be extended to the 1-skeleton, we can assign a sign $+$ or $-$ to each connected component of X and naturally extend to $X \to \{+, -\}$. If $n = 2$, then f' has values in a circle, S^1. Assume that it can be extended to $A \cup X^{(2)} \to S^1$. Then any extension $A \cup X^{(j)} \to S^1$ of f', $j \geq 2$, can be extended to $A \cup X^{(j+1)}$, because the

restriction of g to the boundary of any $(j+1)$-simplex Δ^{j+1} defines a map $\partial\Delta^{j+1} \to S^1$ from a j-sphere to the circle and such map is homotopic to a constant (see, e.g. [36, Chapter 4.1]), hence $\partial\Delta^{j+1} \to S^1$ can always be extended to all of Δ^{j+1}.

Assume that $\dim X \leq n+1$ or $n < 3$ and that $r > r_k$. Then $r > r_0$ and $f'|_{A_r}$ is well defined and homotopic to $f|_{A_r}$ by Lemma 2.4. This implies the extendability of $f|_{A_r}$ to all of X and the relation $A(r+\alpha) \subseteq A_r$ implies the extendability of $f|_{A(r+\alpha)}$. Thus the robustness is less than $r+\alpha$ for any $r > r_k$, yielding an upper bound $r_k + \alpha$ on the robustness of zero. □

To conclude this section, we remark that

- In case when f has no zero at all, we may easily approximate the *robustness of non-existence of zero* by $\min_v |f(v)| - \alpha$.
- The infinite bound in $B(b)$ can be improved to $\max_v |f(v)| + \alpha$.

4 Persistence of obstructions

In this section we describe the algorithm for computing the persistence of primary obstruction and roughly outline the algorithm for secondary obstruction (which is described in more detail in Appendix A).

Primary obstruction—extendability to $X^{(n)}$. Here we review some facts from obstruction theory. A reference for the next proposition can be a textbook such as [42, III, 1.2].

Proposition 4.1 (Primary obstruction)**.** *Let $A \subseteq X$ be a pair of simplicial complexes, $f\colon A \to \Sigma^{n-1}$ be simplicial, $z \in C^{n-1}(\Sigma^{n-1}; \mathbb{Z})$ be a cocycle generating the cohomology, and $y := f^\sharp(z) \in Z^{n-1}(A; \mathbb{Z})$ its pullback.*

Then $f : A \to \Sigma^{n-1}$ can be extended to a (not necessarily simplicial) map $A \cup X^{(n)} \to \Sigma^{n-1}$, iff $y \in Z^{n-1}(A; \mathbb{Z})$ can be extended to a cocycle $x \in Z^{n-1}(X; \mathbb{Z})$ such that $x|_A = y$.

Thus extendability of f is reduced to extendability of an A-cocycle y to a global cocycle x defined on all of X. We will use the notation $\Omega(A) := \{x \in Z^{n-1}(X; \mathbb{Z})\colon x|_A = y\}$ of all cocycle extensions and want to test its non-emptiness. $\Omega(A)$ corresponds to solutions of a linear equation over integers. To see that, let $\bar{y} \in C^{n-1}(X; \mathbb{Z})$ be an arbitrary cochain (not necessarily a cocycle) such that $\bar{y}|_A = y$. We have that

$$\Omega(A) = \{\bar{y} - c\colon c \in C^{n-1}(X, A; \mathbb{Z}) \text{ such that } \delta c = \delta\bar{y}\}. \tag{2}$$

Subtracting $c \in C^{n-1}(X, A; \mathbb{Z})$ does not change the values on A-simplices, so any such $\bar{y} - c$ is still an extension of y. The non-emptiness of $\Omega(A)$ is thus equivalent to solvability of the linear equation $\delta c = \delta \bar{y}$ with the unknown $c \in C^{n-1}(X, A; \mathbb{Z})$.

A natural set of generators of $C^{n-1}(X, A; \mathbb{Z})$ is the set of all $(n-1)$-simplices in X that are not in A with the identification between a simplex Δ and its *characteristic cochain* that assigns 1 to Δ and 0 to all other simplices. Converting $\delta c = \delta \bar{y}$ into an explicit matrix system of linear equations then amounts to enumerating the $(n-1)$- and n-simplices in $X \setminus A$, computing the codifferential matrix of δ using the definition of boundary and expressing the right-hand side $\delta \bar{y}$ in the basis of the n-simplices.

Persistence of the primary obstruction—the algorithm. We remind that in the persistent setting the input contains a filtration of simplicial complexes $X \supseteq A_{r_0} \supseteq A_{r_1}, \ldots, \supseteq A_{r_h}$ and a simplicial map $f' : A_{r_0} \to \Sigma^{n-1}$. We want to compute the smallest value j such that the restriction of f' to A_{r_j} can be extended to $A_{r_j} \cup X^{(n)}$.

Let $A := A_{r_0}$ and z, y be defined as above. The cochain extension $\bar{y}$ of $(f')^\sharp(z)$ is also an extension of $(f'|_{A'})^\sharp(z)$ for each $A' \subseteq A$. Thus we fix one $\bar{y}$ for all spaces A_r. Then the only thing that is changing in solving $\delta c = \delta \bar{y}$, $c \in C^{n-1}(X, A_r; \mathbb{Z})$ with increasing r, is that $X \setminus A_r$ is becoming larger and we are allowed to include more columns into our matrix representing δ.

Now we describe the algorithm on a lower level.

- First we choose the cocycle $z \in \Sigma^{n-1}$ that generates the $(n-1)$ cohomology and its pullback $y := (f'|_A)^\sharp(z) \in Z^{n-1}(A; \mathbb{Z})$ for $A = A_{r_0}$.
- We fix an arbitrary extension $\bar{y} \in C^{n-1}(X; \mathbb{Z})$ of y: the simplest option is to choose $\bar{y}(\Delta) = 0$ for all $(n-1)$-simplices $\Delta \in X$ that are not in A.
- We compute the filtration values of all $(n-1)$ simplices in X. by (1).
- We order the $(n-1)$-simplices of X by their filtration value, and choose an arbitrary enumeration of the n-simplices. These choices will serve as bases of the $(n-1)$- and n-cochains (we identify simplices and their characteristic cochains).
- We construct the matrix M representing the codifferential with respect to the bases chosen above. The columns of M are coboundaries of the $(n-1)$-simplices ordered by increasing filtration values. Further, we convert the right-hand side $\delta \bar{y}$ to an integer vector $\boldsymbol{a}$ using the chosen basis of n-simplices.

Recall that we want to solve $\delta c = \delta\bar{y}$ for c that is a linear combination of $(n-1)$-simplices with filtration values smaller than r_j for j as small as possible. This directly translates to the following problem, which is the last step of the persistence-of-primary-obstruction algorithm.

Problem EARLIEST SOLUTION
Input: A matrix $M \in \mathbb{Z}^{p\times q}$ and a column vector $\boldsymbol{a} \in \mathbb{Z}^p$.
Output: A column vector $\boldsymbol{x} \in \mathbb{Z}^q$ such that $M\boldsymbol{x} = \boldsymbol{a}$.
Objective: Minimize the index of the *lowest nonzero entry* of $\boldsymbol{x}$, that is, $\ell \geq 0$ such that $x_\ell \neq 0$ and $x_{\ell+1} = x_{\ell+2} = \ldots = x_q = 0$.

This problem could be solved by binary search on the value ℓ while solving an ordinary linear system of equations in each iteration. Our implementation uses a simple matrix reduction approach (resembling algorithms for persistent homology) which avoids the binary search (see Appendix B for details).

Secondary obstruction—extendability to $A\cup X^{(n+1)}$. Computing the secondary obstruction and its persistence contains similar ingredients but is more technical and we postpone a lower-level description to Appendix A. Here we outline the main steps for the non-persistent version with a fixed A. We assume that $f' : A \to \Sigma^{n-1}$ is extendable to $A \cup X^{(n)}$ and that $\Omega(A)$ (described by (2)) is nonempty.

We need to implement the "Steenrod square" operation $\smile_{n-3}$: $C^{n-1}(X;\mathbb{Z}_2)\times C^{n-1}(X;\mathbb{Z}_2) \to C^{n+1}(X;\mathbb{Z}_2)$ on the level of cocycles. The algorithm for $\smile_{n-3}$ directly follows from formulas in [44, p. 292–293]. For the following facts, we refer to [44] and [45]:

Proposition 4.2 (Secondary obstruction)**.** *Let $A \subseteq X$ be a pair of simplicial complexes, $f' : A \to \Sigma^{n-1}$ be simplicial and assume that the ordering of vertices of X and Σ^{n-1} is chosen so that $v \preceq w \;\Rightarrow\; f'(v) \preceq f'(w)$. For each $x \in \Omega(A)$ let $(x \mod 2)$ be the image of x under the natural homomorphism $C^{n-1}(X;\mathbb{Z}) \to C^{n-1}(X;\mathbb{Z}_2)$. Then*

$$v(x) := (x \mod 2) \smile_{n-3} (x \mod 2) \tag{3}$$

vanishes on A, that is, it is an element of $Z^{n+1}(X, A;\mathbb{Z}_2)$.

Further, if $n > 3$, then f' can be extended to a map $X^{(n+1)} \to \Sigma^{n-1}$ iff $v(x)$ is a relative coboundary for some $x \in \Omega(A)$.

Thus extendability to $X^{(n+1)}$ is equivalent to satisfiability of the equation $\delta c = v(x)$, $c \in C^n(X, A;\mathbb{Z}_2)$, for *some* $x \in \Omega(A)$. To decide this, we parameterize $\Omega(A)$ by a fixed representant x and generators g_j of $Z^{n-1}(X, A;\mathbb{Z})$:

arbitrary element of $\Omega(A)$ is then $x - \sum_j u_j g_j$ for some $u_j \in \mathbb{Z}$. To reduce the number of j's, we only need to take generators of the cohomology group $H^{n-1}(X, A; \mathbb{Z})$. Exploiting the "linearity" of the operation v on the level of cohomology ([44, p. 504]), we have that $v(x - \sum_j u_j g_j)$ is a coboundary iff $v(x) - \sum_j u_j v(g_j)$ is a coboundary. Thus our equation reduces to $\delta c + \sum_j u_j v(g_j) = v(x)$. This is a system of equations with right-hand side $v(x)$ and unknowns c and u_j, this time over the $\mathbb{Z}_2$-coefficients.

We also remark that the last proposition is valid also in the case $n = 3$ once we replace $\mathbb{Z}_2$-coefficients by $\mathbb{Z}$-coefficients, but deciding satisfiability of $\delta c = v(x)$ for some x is hard (and undecidable for general spaces X). We show at the end of Appendix A that if X is a triangulation of the topological cube $[0, 1]^4$ and $n = 3$, then triviality of the secondary obstruction can easily be tested as well: this case is also included in our implementation.

To compute the smallest r such that the map $f'|_{A_r}$ is extendable to $A_r \cup X^{(n+1)}$, we could use a binary search. As in the case of the primary obstruction, it can be avoided and we can compute the persistence of the secondary obstruction using a single matrix reduction: this is explained in Appendix A.

5 Application for robust optimization

Reduction of robust optimization to the ROB-SAT problem. Our algorithmic approach has a natural extension for optimization with uncertainty. We pose the following optimization problem:

$$\begin{array}{ll} \text{maximize} & o(x) \\ \text{subject to} & g(x) = 0 \\ & x \in X \end{array} \tag{4}$$

where X is a compact domain[5] and $g\colon X \to \mathbb{R}^n$ is uncertain. Let as assume, for simplicity, that $r > 0$ is fixed and g is an unknown continuous r-perturbation of a known given map $f\colon X \to \mathbb{R}^n$. A simple instance of the above problem is visualized below:

[5]Such domain implicitly imposes inequality constraints which can be seen as uncertain ones as well if the chosen norm on $\mathbb{R}^n$ is ℓ_∞, see [26]. Also the function o could be considered as uncertain without adding further complexity to the problem, but we prefer to have the statement as simple as possible.

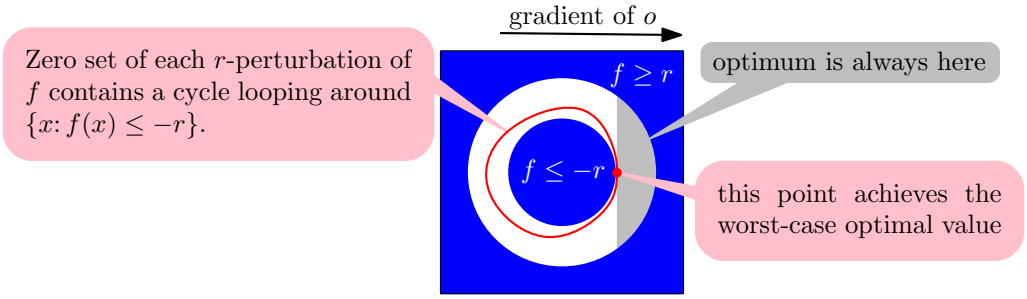
gradient of o
Zero set of each r-perturbation of f contains a cycle looping around $\{x\colon f(x) \le -r\}$.
$f \ge r$
$f \le -r$
optimum is always here
this point achieves the worst-case optimal value

Both f and $o\colon X \to \mathbb{R}$ can be specified in various ways but let as further assume that they are *simplexwise linear* and that we know their values on vertices.

We remark that in the mainstream approach the uncertainty is parameterized, that is, in the problem above, the constraints would have the form $g_p(x) = 0$ where p is an unknown vector-valued parameter.

The common goal is to compute the optimal value in the worst case, i.e.,

$$\inf_{\|g-f\|\le r} \max_{x\in g^{-1}(0)} o(x)$$

in our case. The worst-case optimal value is equal to the maximal number $\beta \in \mathbb{R}$ such that f has an r-robust zero on $o^{-1}[\beta,\infty)$. If o is simplexwise linear, then $o^{-1}[\beta,\infty)$ can be triangulated and the existence of an r-robust zero on $o^{-1}[\beta,\infty)$ can be algorithmically tested via computing higher-order obstructions, whenever $\dim X \le 2n-3$ or $n < 3$ [26]. The exact worst-case optimal value can be found by doing a binary search on the maximal value β and using the ROB-SAT algorithm [26] in each step.

Efficient implementation. Also the efficiency–tuned algorithm presented in this paper can be tweaked into the setting of optimization very easily and thus the binary search avoided. We may assume that both f and o are only given via function values in vertices and simplex-wise Lipschitz constants, and want to approximate $\inf_{\|g-f\|\le r} \max_{x\in g^{-1}(0)} o(x)$ for some r.[6]

The only difference occurs before each call of Earliest Solution subroutine where we sort the rows of the matrix M and the right-hand side $\boldsymbol{a}$ (n-simplices in the case of primary obstruction) according to their o-filtration value (minimum of $o(v)$ over their vertices v). Also we cut off the columns of the matrix with filtration value larger than r. After the column matrix

[6] To avoid further simplicial subdivisions, we again need to assume that $r > \alpha n^{1/p}$, i.e., that the description of f is fine-grained enough for the retrieval of the homotopy class of $f|_{A_r}$).

reduction as described in Appendix B, the desired approximation of the worst-case optimal value is the o-filtration value corresponding to the row of the lowest nonzero element on the right hand side $\boldsymbol{a}$ after the reduction.

We remark that we can immediately compute a lower bound[7] on the *uncertainty-optimality curve* $\mathrm{OPT}(r) := \inf_{\|g-f\|\leq r} \max_{x\in g^{-1}(0)} o(x)$ as the o-filtration value of the lowest nonzero entry of the right-hand side after the reduction by the column of filtration value r. These values are just a side product of the matrix reduction algorithm in Appendix B. The error in this approximation is bounded by the simplexwise Lischitz constants for f and o.

6 Experimental results

Motivation. One motivation for implementing the algorithm was to experimentally analyse the following question:

- How typical is a situation in which the zero cannot be detected by primary obstruction and higher obstructions are needed?

To illustrate the flavour of this problem, consider a function f from an $(n+1)$-ball B^{n+1} to $\mathbb{R}^n$ such that the zero set is a circle. If r is small, then the r-neighborhood of the zero set is homeomorphic to a solid torus $S^1 \times B^n$. An n-hyperplane intersecting the zero set transversally will typically intersect this torus in a n-disc $\{*\} \times B^n$ with a zero of f inside: this reflects the non-extendability to the n-skeleton. However, with increasing r (and hence increasing our freedom to perturbate the function), the primary obstruction will die once the r-neighborhood will touch the boundary or become a full $(n+1)$-ball: in the latter case, a nontrivial secondary obstruction is reflected by the homotopy class of the map from the boundary of this $(n+1)$-ball to S^{n-1}. This homotopy class is encoded in the gradient-induced framing of the original zero set of f: if the framing is trivial (framed null-cobordant), then higher obstruction don't occur. If the framing is "twisted", then they do.

Intuitively, we assumed that using Gaussian random fields, the gradient-induced framing of the zero set should be quite random and we would observe twisted as well as untwisted cases. Experiments, however, do not support this so far, which we find surprising.

[7] An upper bound is obtained when the dimension is at most n or $n \leq 2$ for the primary obstruction and at most $n+1$ for the secondary obstruction.

Description of the computation experiments. The lowest-dimensional case where nontrivial secondary obstruction can occur is $\dim X = 4$ and $n = 3$. Using an experimental approach, we generated random continuous functions from a regular 4-dimensional cubical grid into $\mathbb{R}^3$ taken from different probability distributions. The space X was either a 4-cube or a 4-torus $(S^1)^4$ and the underlying simplicial complex was the Freudenthal triangulation of the canonical cubical subdivision of X [5, p. 154]. Instead of A_r from Definition 2.2, we used a courser filtration $A^{\square}_r$ based on the cubical structure, see Appendix C for details. We computed the vertex-approximation f' from Definition 2.2 and the smallest $r_0 > 0$ such that f' is simplicial on $A^{\square}_{r_0}$. Then we found the persistence of the primary obstruction $r_1 \geq r_0$ and the persistence of the secondary obstruction $r_2 \geq r_1$: the goal was to check whether instances with $r_2 > r_1$ occur and how often.

First we experimented with Gaussian random fields. Such functions are continuous and infinitely differentiable [1, Sec. 2.2]. For each component f_i of f and each vertex x, the random variable $f_i(x)$ was normalized to the standard normal distribution $N(0, 1)$ and the covariance between $f_i(x)$ and $f_i(y)$ was taken to be $C(x, y) = \tilde{C}(|x - y|)$: we tried different functions C. First we generated random functions such that the discrete Fourier transform of $C(0, x)$ was proportional to $((1 + |p|^2)^{-l})_{p\in\{0,\ldots,g-1\}^4}$ for various constants l (compare [39, p. 12]). The value $l = 0$ corresponds to white noise and $l = \infty$ to constant functions. While this procedure naturally creates functions on a torus, for experiments on a cube we generated a random function on the discrete torus $\{1, 2, \ldots, 2g\}^4$ and restricted it to the coordinates $\{1, \ldots, g\}^4$ to avoid periodicity. The three components of f were generated independently. To assure that the resulting function has zero at all, we analyzed the function $f(x) - f(x_0)$ instead of $f(x)$, where x_0 was the midpoint of the cube, resp. a fixed point in the torus.

In most cases, we detected a nontrivial primary obstruction, but not a single instance with secondary obstruction $r_2 > r_1$. To give an illustration, the following table shows some statistics of one of the experiments on a 4-cube: l is the parameter of the distribution, g is the number of vertices in each dimension, r_0 the smallest value for which $f'|_{A^{\square}_{r_0}}$ is simplicial, r_1 the average persistence of the primary obstruction in cases when $r_1 > r_0$, and max. r_1 the largest persistence of primary obstruction. The averages are taken out of 1000 functions for $l \in \{3, 3.5, 4, 4.5\}$ and out of 10 000 for $l = 5.0$.

l	g	r_0	% of $r_1 > r_0$	average r_1 if nontrivial	max. r_1
3.0	30	0.4	78%	0.56	0.95
3.5	30	0.21	91%	0.41	0.82
4.0	25	0.15	91%	0.32	0.66
4.5	25	0.1	91%	0.25	0.55
5.0	20	0.1	87%	0.21	0.63

When performing such experiments on the 4-torus, it sometimes happened that the cup square of a computed extension $x \in \Omega(A_{r_1})$ was nontrivial in $H^4(X, A^{\square}_{r_1})$, giving some "hope" of a nontrivial secondary obstruction: however, in all cases, this could be removed after replacing x by another extension of the pullback $y = (f'|_A)^*(z)$ to the 2-skeleton (see Section 4).[8]

In other rounds of experiments, we generated functions from a 5-torus into $\mathbb{R}^4$ or replaced the correlation function $C(x, y)$ by the Gaussian function $\exp\left(-\frac{|x-y|^2}{2l^2}\right)$ for suitable $l > 0$, but the results were were similar to that from the distribution above.

One possible explanation for the lack of secondary obstruction is that the cohomology in dimension 4 has typically lower persistence than in dimension 3 and most generators have already died when the primary obstruction (element of H^3) dies. A similar phenomenon occurs in persistent homology of excursion sets of random scalar fields, where the persistence barcodes in dimension 0 die before the barcodes in dimension 1, compare [2]. In the vast majority of our experiments on the cube, the 4-dimensional cohomology group $H^4(X, A^{\square}_{r_1})$ was trivial for r_1 being the persistence of primary obstruction. The lack of top dimensional cohomology in this case probably reflects the fact that most components of the neighborhood of the zero set intersect the boundary of the domain, although this argument does not apply for the torus.

In another attempt to detect secondary obstruction in random fields we generated random homogenous quadratic polynomials on $[-1, 1]^4$. The coefficients a^k_{ij} in $f_k(x) = \sum_{i,j} a^k_{i,j} x_i x_j$ were independent samples from a standard normal distribution.[9] The zero set of homogenous quadratic functions is either the origin alone or a cone intersecting the boundary $\partial[-1, 1]^4$: only the first case can yield a nontrivial $H^4(X, A^{\square}_r)$ and a nontrivial sec-

[8] In fact, nontriviality of the secondary obstruction on a 4-torus can only be reduced to a system of quadratic Diophantine equations. While we cannot algorithmically check satisfiability of quadratic equations, in all cases where we had to deal with this problem, these equations were almost trivial and solvable.

[9] This is motivated by the fact that the simplest examples of functions with nontrivial secondary obstruction are quadratic and homogenous.

ondary obstruction. We generated around 70 thousand instances of random quadratic functions on a 10^4 grid: around 2.2% of them had only the origin as the zero set, but there was no nontrivial secondary obstruction in a single instance.

Experiments with formulas. Another motivation for implementing the algorithm was to test the running time and memory limitations in practice. Our testing benchmarks consisted of inputs in which the function values $f(v)$ were generated via formulas with known properties in a cubical grid. We ran many testing examples, some of them being shown in Appendix D. To summarize the results, the performance is much better than the worst-case complexity bound derived in Section 7 and is approximately linear in the number of simplices of the input. We were able to run benchmarks up to $\dim X = 8$ for small grids, such as 5^8: the largest coboundary matrix for which we computed a nontrivial obstruction had 40 million columns.

7 Complexity

The input size (and hence computational complexity) depends heavily on the encoding of the simplicial complex. For example, we may specify the set of *all* simplices, or the set of all top-dimensional simplices.[10] Therefore we study *parameterized complexity* as a function of the following parameters. Let m be the dimension of X and $n \leq m$ the dimension of the target space $\mathbb{R}^n$. We define N to be the maximum of the number of k-simplices for $k \in \{n-1, n, n+1\}$ and V the number of vertices. In addition to specifying X, the input contains the function values in all vertices, that is, $n \times V$ numbers.

We present complexity bounds as a function of N, V, m, n.

Primary obstruction—complexity. We assume that the function values $f(v)$ at the vertices are all rational vectors and that we can compare their absolute values $|f(v)|$, $|f(w)|$ in unit time (these numbers may be roots of rational numbers for ℓ_p norms). Then computing the vertex approximation $f'(v)$ for each vertex v via (2.3) amounts to $O(nV)$ operations. Computing the filtration of all $(n-1)$-simplices via formula (1), as well as the pullback y and its codifferential are by definition subroutines of complexity $O(nN)$; ordering the $(n-1)$-simplices by filtration is done in $O(N \log N)$. The computation of the codifferential matrix is again of order $O(nN)$ if we store

[10] In other situations, the input specifying the simplicial complex could be even smaller. One example is specifying the vertex set in $\mathbb{R}^m$ and assuming the Delaunay triangulation.

it in a sparse format, because each row of the matrix corresponds to the boundary of an n-simplex and has only $n+1$ nonzero elements.

The bottleneck of computing the primary obstruction is the EARLIEST SOLUTION algorithm described on page 13. An implementation based on a binary or exponential search requires at most $\log N$ solutions of a linear system of Diophantine equations. Each of them is a system of at most N rows and columns, all coefficients being ± 1 or 0. By [46, Thm. 19] we may solve such Diophantine system in $O(N^4 \log^4 N)$ time, which yields $O(N^4 \log^5 N + n(N+V))$ as a complexity bound for the primary obstruction. Assuming the lack of blowup of matrix coefficients during the matrix reduction, we can bound the number of arithmetic operations in EARLIEST SOLUTION by $O(N^3)$. This is discussed in more detail in Appendix B. In this scenario, sub-cubic bounds could be achieved using randomization [47, Thm. 39]. In practice, however, our implementation of EARLIEST SOLUTION exhibits subquadratic scaling, allowing us to experiment with instances for $N \leq 10^7$. This is not entirely surprising—large instances of simplicial boundary matrices are commonly reduced in the field of computational topology.

Secondary obstruction—complexity. The bottleneck in the secondary obstruction algorithm is the computation of all Steenrod squares of all generators of $H^{n-1}(X, A_r; \mathbb{Z})$ for all filtration values r. In a naive implementation we may compute a set of generators of $Z^{n-1}(X; \mathbb{Z})$ and their respective filtration values. Generators of the kernel (over $\mathbb{Z}$) of a matrix with at most N rows and columns can be computed in $O(N^4)$ [16, Theorem 1]. The number of such generators is bounded by N. In the Steenrod square computation, we need to compute, in the worst case, the values on all $(n+1)$-simplices; in each evaluation, the formula for $\smile_{n-3}$ described in [44] contains an iteration over all elements of $\binom{m}{4}$ (Steenrod pairs). Thus, computing the Steenrod squares of the generators of $Z^{n-1}(X; \mathbb{Z})$ is $O(N^2\, m^4)$. The final matrix computation corresponding to equation (7) is done over the field $\mathbb{Z}_2$ which only requires $O(N^\omega)$ operations for a constant $\omega < 3$ [37, Proposition 6]. This yields a complexity bound of $O(N^4 + N^2 m^4)$ for the persistence of secondary obstruction. For all practical purposes, it is safe to assume that the values of m can be neglected.

References

[1] Adler, R.J.: The geometry of random fields, vol. 62. Siam (1981)

[2] Adler, R.J., Bobrowski, O., Borman, M.S., Subag, E., Weinberger, S.: Persistent homology for random fields and complexes (2010)

[3] Alefeld, G., Frommer, A., Heindl, G., Mayer, J.: On the existence theorems of Kantorovich, Miranda and Borsuk. Electronic Transactions on Numerical Analysis **17**, 102–111 (2004)

[4] Alefeld, G.E., Shen, Z.: Miranda's theorem and the verification of solution of linear complementarity problems. Tech. Rep. 01/05, Institut für Wissenschaftliches Rechnen und Mathematische Modellbildung (2001)

[5] Allgower, E.L., Georg, K.: Introduction to numerical continuation methods, vol. 45. SIAM (2003)

[6] Aubry, C., Desmare, R., Jaulin, L.: Loop detection of mobile robots using interval analysis. Automatica **49**(2), 463 – 470 (2013). DOI http://dx.doi.org/10.1016/j.automatica.2012.11.009. URL http://www.sciencedirect.com/science/article/pii/S0005109812005456

[7] Aubry, C., Desmare, R., Jaulin, L.: Kernel characterization of an interval function. Mathematics in Computer Science **8**(3), 379–390 (2014). DOI 10.1007/s11786-014-0206-9. URL http://dx.doi.org/10.1007/s11786-014-0206-9

[8] Bauer, U., Kerber, M., Reininghaus, J., Wagner, H.: Phat–persistent homology algorithms toolbox. In: Mathematical Software–ICMS 2014, pp. 137–143. Springer Berlin Heidelberg (2014)

[9] Ben-Tal, A., Ghaoui, L., Nemirovski, A.: Robust Optimization. Princeton Series in Applied Mathematics. Princeton University Press (2009). URL http://books.google.cz/books?id=DttjR7IpjUEC

[10] Ben-Tal, A., Nemirovski, A.: Robust optimization – methodology and applications. Mathematical Programming **92**(3), 453–480 (2002). DOI 10.1007/s101070100286. URL http://dx.doi.org/10.1007/s101070100286

[11] Bendich, P., Edelsbrunner, H., Morozov, D., Patel, A.: The robustness of level sets. In: M. Berg, U. Meyer (eds.) Algorithms–ESA 2010, *Lecture Notes in Computer Science*, vol. 6346, pp. 1–10. Springer (2010). DOI 10.1007/978-3-642-15775-2_1. URL http://dx.doi.org/10.1007/978-3-642-15775-2_1

[12] Bendich, P., Edelsbrunner, H., Morozov, D., Patel, A.: Homology and robustness of level and interlevel sets. Homology, Homotopy and Applications **15**(1), 51–72 (2013). URL http://projecteuclid.org/euclid.hha/1383943667

[13] Bertsimas, D., Brown, D.B., Caramanis, C.: Theory and applications of robust optimization. SIAM review **53**(3), 464–501 (2011)

[14] Beyer, H.G., Sendhoff, B.: Robust optimization – a comprehensive survey. Computer Methods in Applied Mechanics and Engineering **196**(33–34), 3190 – 3218 (2007). DOI http://dx.doi.org/10.1016/j.cma.2007.03.003. URL http://www.sciencedirect.com/science/article/pii/S0045782507001259

[15] Bredon, G.: Topology and Geometry. Graduate Texts in Mathematics 139. Springer-Verlag, Berlin etc. (1993)

[16] Buchmann, J., Squirrel, D.: Kernels of integer matrices via modular arithmetic. Tech. rep. (1999). URL https://www.researchgate.net/publication/2611992_Kernels_of_Integer_Matrices_via_Modular_Arithmetic

[17] Čadek, M., Krčál, M., Matoušek, J., Sergeraert, F., Vokřínek, L., Wagner, U.: Computing all maps into a sphere. J. ACM **61**(3), 17:1–17:44 (2014). DOI 10.1145/2597629. URL http://doi.acm.org/10.1145/2597629

[18] Čadek, M., Krčál, M., Matoušek, J., Vokřínek, L., Wagner, U.: Extendability of continuous maps is undecidable. Discr. Comput. Geom. **51**(1), 24–66 (2013). To appear. Preprint arXiv:1302.2370

[19] Čadek, M., Krčál, M., Matoušek, J., Vokřínek, L., Wagner, U.: Polynomial-time computation of homotopy groups and Postnikov systems in fixed dimension. Siam Journal on Computing **43**(5), 1728–1780 (2014)

[20] Chazal, F., Patel, A., Škraba, P.: Computing the Robustness of Roots. Applied Mathematics Letters **25**(11), 1725 — 1728 (2012). URL http://ailab.ijs.si/primoz_skraba/papers/fp.pdf

[21] Chung, M.K., Bubenik, P., Kim, P.T.: Information Processing in Medical Imaging: 21st International Conference, IPMI 2009, Williamsburg, VA, USA, July 5-10, 2009. Proceedings, chap. Persistence Diagrams

of Cortical Surface Data, pp. 386–397. Springer Berlin Heidelberg, Berlin, Heidelberg (2009). DOI 10.1007/978-3-642-02498-6_32. URL http://dx.doi.org/10.1007/978-3-642-02498-6_32

[22] Dian, J., Kearfott, R.B.: Existence verification for singular and nonsmooth zeros of real nonlinear systems. Mathematics of Computation **72**(242), 757–766 (2003)

[23] Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. Discrete and Computational Geometry **28**(4), 511–533 (2002)

[24] Eilenberg, S., Zilber, J.A.: On products of complexes. American Journal of Mathematics pp. 200–204 (1953)

[25] Franek, P., Krčál, M.: On Computability and Triviality of Well Groups. In: L. Arge, J. Pach (eds.) 31st International Symposium on Computational Geometry (SoCG 2015), *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 34, pp. 842–856. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2015). DOI http://dx.doi.org/10.4230/LIPIcs.SOCG.2015.842

[26] Franek, P., Krčál, M.: Robust satisfiability of systems of equations. J. ACM **62**(4), 26:1–26:19 (2015). DOI 10.1145/2751524. URL http://doi.acm.org/10.1145/2751524

[27] Franek, P., Ratschan, S.: Effective topological degree computation based on interval arithmetic. AMS Mathematics of Computation **84**(293), 1265–1290 (2015)

[28] Franek, P., Ratschan, S., Zgliczynski, P.: Satisfiability of systems of equations of real analytic functions is quasi-decidable. In: *Proc. 36th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, LNCS 6907, pp. 315–326. Springer, Berlin (2011)

[29] Franek, P., Ratschan, S., Zgliczynski, P.: Quasi-decidability of a fragment of the first-order theory of real numbers. Journal of Automated Reasoning pp. 1–29 (2015). DOI 10.1007/s10817-015-9351-3. URL http://dx.doi.org/10.1007/s10817-015-9351-3

[30] Friedman, G.: An elementary illustrated introduction to simplicial sets. Rocky Mountain J. Math. **42**(2), 353–423 (2012)

[31] Frommer, A., Lang, B.: Existence tests for solutions of nonlinear equations using Borsuk's theorem. SIAM Journal on Numerical Analysis **43**(3), 1348–1361 (2005). DOI 10.1137/S0036142903438148. URL http://link.aip.org/link/?SNA/43/1348/1

[32] Gao, M., Chen, C., Zhang, S., Qian, Z., Metaxas, D., Axel, L.: Segmenting the papillary muscles and the trabeculae from high resolution cardiac ct through restoration of topological handles. In: International Conference on Information Processing in Medical Imaging (IPMI) (2013)

[33] Goldsztejn, A., Jaulin, L.: Inner approximation of the range of vector-valued functions. Reliable Computing pp. 1–23 (2010)

[34] Goldsztejn, R., Jaulin, L.: Inner and outer approximations of existentially quantified equality constraints. In: In Proceedings of the Twelfth International Conference on Principles and Practice of Constraint Programming, (CP 2006 (2006)

[35] Gonzalez-Diaz, R., Real, P.: Simplification techniques for maps in simplicial topology. Journal of Symbolic Computation **40**(4), 1208–1224 (2005)

[36] Hatcher, A.: Algebraic Topology. Cambridge University Press, Cambridge (2001). URL https://www.math.cornell.edu/~hatcher/AT/ATpage.html

[37] Jeannerod, C.P., Pernet, C., Storjohann, A.: Rank-profile revealing gaussian elimination and the cup matrix decomposition. Journal of Symbolic Computation **56**, 46–68 (2013)

[38] Krčál, M., Matoušek, J., Sergeraert, F.: Polynomial-time homology for simplicial Eilenberg–MacLane spaces. J. Foundat. of Comput. Mathematics **13**, 935–963 (2013). Preprint, arXiv:1201.6222

[39] Lang, A., Potthoff, J.: Fast simulation of gaussian random fields. Monte Carlo Methods and Applications **17**(3), 195–214 (2011)

[40] Maria, C., Boissonnat, J.D., Glisse, M., Yvinec, M.: The gudhi library: Simplicial complexes and persistent homology. In: H. Hong, C. Yap (eds.) Mathematical Software – ICMS 2014, *Lecture Notes in Computer Science*, vol. 8592, pp. 167–174. Springer Berlin Heidelberg (2014). DOI 10.1007/978-3-662-44199-2_28. URL http://dx.doi.org/10.1007/978-3-662-44199-2_28

[41] Merlet, J.P.: Interval analysis and reliability in robotics. International Journal of Reliability and Safety **3**(1-3), 104–130 (2009)

[42] Prasolov, V.V.: Elements of Homology Theory. Graduate Studies in Mathematics. American Mathematical Society (2007)

[43] Rump, S.M.: Verification methods: Rigorous results using floating-point arithmetic. Acta Numerica **19**, 287–449 (2010). DOI 10.1017/S096249291000005X. URL http://journals.cambridge.org/article_S096249291000005X

[44] Steenrod, N.E.: Products of cocycles and extensions of mappings. Annals of Mathematics **48**(2), pp. 290–320 (1947)

[45] Steenrod, N.E.: Cohomology operations, and obstructions to extending continuous functions. Advances in Math. **8**, 371–416 (1972)

[46] Storjohann, A.: A fast+ practical+ deterministic algorithm for triangularizing integer matrices (1996). URL http://e-collection.library.ethz.ch/eserv/eth:3348/eth-3348-01.pdf

[47] Storjohann, A.: The shifted number system for fast linear algebra on integer matrices. Journal of Complexity **21**(4), 609–650 (2005)

[48] Vokřínek, L.: Decidability of the extension problem for maps into odd-dimensional spheres. ArXiv e-prints (2014)

[49] Wang, P.S.: The undecidability of the existence of zeros of real elementary functions. J. ACM **21**(4), 586–589 (1974). DOI http://doi.acm.org/10.1145/321850.321856

A Secondary obstruction.

Persistence of the secondary obstruction—the algorithm for $n > 3$. Assume that a filtration $X \supseteq A_{r_0} \supseteq A_{r_1} \supseteq \dots$ and a simplicial map $f' : A_{r_0} \to \Sigma^{n-1}$ are given, $n > 3$, and vertices on X and Σ^{n-1} are ordered so that f' is order-preserving: this order is used in the implementation of the $\smile_{n-3}$ operation on the level of cochains. Further, we assume that the persistence of primary obstruction r_j has already been computed by the algorithm described on page 12. We continue to use the notation of Section 4: in particular, z is the characteristic cocycle of a fixed $(n-1)$-simplex in Σ^{n-1}, $\bar{y} \in C^{n-1}(X, A_{r_j})$ is a cochain extending the pullback $y = (f')^\sharp(z) \in Z^{n-1}(A_{r_j})$ of z and $\emptyset \neq \Omega(A_{r_j})$ is the set of all $(n-1)$-cocycles on X that extend y on A_{r_j}.

By Proposition 4.2, the persistence of secondary obstruction is the smallest number r_k such that

$$\delta c = v(x) \quad \text{for some } c \in C^n(X, A_{r_k}; \mathbb{Z}_2), x \in \Omega(A_{r_k}) \tag{5}$$

where v is defined by (3).

Let $x \in \Omega(A_{r_j})$ be a fixed extension of $\bar{y}$, computed in the algorithm for primary persistence. Then also $x \in \Omega(A_{r_k})$ for each $k \geq j$. For any k, $\Omega(A_{r_k})$ is a coset in $Z^{n-1}(X)$

$$\Omega(A_{r_k}) = x + Z^{n-1}(X, A_{r_k}; \mathbb{Z})$$

and hence equation (5) reduces to

$$\delta c = v(x - w), \quad \text{for some } w \in Z^{n-1}(X, A_{r_k}; \mathbb{Z}),\ c \in C^n(X, A_{r_k}; \mathbb{Z}_2). \tag{6}$$

The crucial property that we will use is that $v(x-w)$ is a relative coboundary iff $v(x)-v(w)$ is a relative coboundary: this follows directly from the linearity of the Steenrod square operation $H^{n-1}(X, A; \mathbb{Z}_2) \to H^{n+1}(X, A; \mathbb{Z}_2)$ for $n > 3$. Thus we can reformulate (6) to the problem of finding the minimal r_k such that

$$\delta c + v(w) = v(x) \quad \text{for some } w \in Z^{n-1}(X, A_{r_k}; \mathbb{Z}),\ c \in C^n(X, A_{r_k}; \mathbb{Z}_2). \tag{7}$$

To simplify the computations, we don't need to consider all cocycles $w \in Z^{n-1}(X, A_r)$ but only generators of the cohomology $H^{n-1}(X, A_r; \mathbb{Z})$: the Steenrod square of any relative coboundary is again a relative coboundary $\delta c'$ for some $c' \in C^n(X, A_r; \mathbb{Z}_2)$, so adding it has no impact on the solvability of (7).

The right-hand side of (7) $v(x)$ is a cocycle that does not depend on k. The left-hand side is a combination of coboundaries of characteristic cocycles of n-simplices Δ^n and cochains of the form $v(w)$ for $(n-1)$-cocycles w. To each Δ and w is assigned a filtration value and we want to minimize the maximal filtration r_k of all the Δ^n's and w's.

Summarizing the above steps, we obtain the following algorithm:

- Order the vertices of Σ^{n-1} and the vertices of X so that f is order-preserving
- For a precomputed $x \in \Omega(A_{r_j})$, compute the relative cocycle $v(x) \in Z^{n+1}(X, A_{r_j}; \mathbb{Z}_2)$ by the definition of the Steenrod operation $\smile_{n-3}$
- Compute a subset $W \subseteq Z^{n-1}(X)$ that contains all cohomology generators w of all $H^{n-1}(X, A_r; \mathbb{Z})$ for all $r \geq r_j$. (It may be a set of generators of $Z^{n-1}(X)$.) To each $w \in W$, assign its filtration value $r(w)$ to be the minimal r such that w is zero on A_r.
- Compute the filtration value of all n-simplices, using (1). Order
- Order all n-simplices and elements of W by their filtration.
- Choose a basis of $(n+1)$-simplices and express the right-hand side $v(x)$ as a vector $\boldsymbol{a} \in (\mathbb{Z}_2)^q$.
- Create a matrix M whose column set consists of
  - all coboundaries of characteristic cochains of n-simplices expressed in the basis of $(n+1)$-simplices over $\mathbb{Z}_2$
  - all elements $v(w)$ expressed in the basis of $(n+1)$-simplices
- Order the columns of M by filtration.

Computing the persistence of the secondary obstruction then reduces to solving the EARLIEST SOLUTION problem for $M\boldsymbol{x} = \boldsymbol{a}$, this time over the $\mathbb{Z}_2$-field.

The hardest part is to compute the cohomology generators: this algorithm is summarized as follows.

Problem PERSISTENT GENERATORS:
Input: A filtration $X \supseteq A_0 \supseteq \ldots \supseteq A_h = \emptyset$.
Output: A sequence $g_1, \ldots, g_\nu \in Z^{n-1}(X; \mathbb{Z})$ and a sequence of integers $\mu(0) \leq \ldots \leq \mu(h)$ such that $[g_1], \ldots, [g_{\mu(i)}]$ generate $H^{n-1}(X, A_i; \mathbb{Z})$ for each i.

We give imeplementation details of this part on a lower-level in Section B.2. Our algorithm will give the output with the number of generators ν minimal in a certain sense. Notably, when working over $\mathbb{Q}$ or $\mathbb{Z}_p$ instead of integers, the output would correspond to a persistence bar code with all the death information erased but including representative (co)cycles for each bar.

The special case $f : [0,1]^4 \to \mathbb{R}^3$. This section justifies the claim that for $X = [0,1]^4$ and $n = 3$, the general algorithm for computing persistence of secondary obstruction works, if we completely ignore the persistence generators w and perform all computations over $\mathbb{Z}$-coefficients. The following text is not completely self-contained, as the proof of Lemma A.1 requires some definitions and techniques from the field of algebraic topology.

Assume first that $n = 3$ and X is arbitrary. The two main differences compared to the algorithm above are:

- The Steenrod operation and final coboundary matrix have to be computed over $\mathbb{Z}$, not $\mathbb{Z}_2$. This goes back to the fact that the homotopy group $\pi_3(S^2) \simeq \mathbb{Z}$, unlike $\pi_n(S^{n-1}) \simeq \mathbb{Z}_2$ for $n > 3$. The homotopy group serves as cohomology coefficients in the theory of obstructions.
- The operation $\smile_{n-3}$ reduces to the cup product $\smile$ and the operation $w \mapsto w \smile w$ is not linear on the level of cohomology but quadratic.

This means that while for any particular extension $x \in \Omega(A)$ of $\bar{y}$, we may test satisfiability of $\delta c = v(x)$, we cannot test the *existence* of such x using a linear system of equations. However, if $X = [0,1]^4$, we claim that for any two extensions $x, y \in \Omega(A)$ of $\bar{y}$, $v(x) - v(y)$ is a relative coboundary and thus we need to check the equation $\delta c = v(x)$ only for one x.

Lemma A.1. *Let X be a triangulation of $[0,1]^4$, $A \subseteq X$, $x \in Z^2(X)$ and $w \in Z^2(X, A)$. Then*

$$(x - w) \smile (x - w) - (x \smile x) \in B^4(X, A).$$

Using the parametrization $\Omega(A) = \{x - w : \ w \in Z^{n-1}(X, A; \mathbb{Z})\}$ we immediately obtain that our general algorithm for the persistence of secondary obstruction works once we replace the $\mathbb{Z}_2$-coefficients by $\mathbb{Z}$-coefficients in its final step. Moreover, we may completely ignore the persistent generators w and don't need to compute them at all.[11]

[11] Note that $x \smile x$ may represent a nontrivial element of $H^4(X, A)$, as x is not an element of $Z^2(X, A)$ in general. The fact that $x \smile x$ is zero on A follows from the fact that f is order-preserving.

Proof. By bilinearity of the cup product, $(x-w) \smile (x-w) - x \smile x = -x \smile w - w \smile x + w \smile w$. The mixed-term $x \smile w$ is a relative coboundary, because $\smile$ induces a bilinear product on the level of cohomology $H^2(X) \times H^2(X, A) \to H^4(X, A)$ and $H^2(X) = 0$, as X is contractible. It remains to show that $w \smile w$ is a coboundary. Let CA be the cone over A and $A \hookrightarrow CA$ be an inclusion. The inclusion $A \hookrightarrow X$ can be extended to a map $CA \to X$, because X is contractible. The map of pairs $(CA, A) \to (X, A)$ induces the commutative diagram in which the rows are the long exact sequences of cohomology groups.

$$\begin{array}{ccccccccc} H^{*-1}(X) & \to & H^{*-1}(A) & \to & H^*(X,A) & \to & H^*(X) & \to & H^*(A) \\ \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\ H^{*-1}(CA) & \to & H^{*-1}(A) & \to & H^*(CA,A) & \to & H^*(CA) & \to & H^*(A) \end{array}$$

The vertical arrows $H^{*(-1)}(X) \to H^{*(-1)}(CA)$ are trivial as both spaces are contractible and $H^{*(-1)}(A) \to H^{*(-1)}(A)$ are identities. By the five-lemma [36, p. 129], the middle homomorphism $H^*(X, A) \to H^*(CA, A)$ is an isomorphism. Further, $H^j(CA, A) \simeq H^j(CA/A)$ for $j > 0$. The space $CA/A =: \Sigma A$ is the *suspension* of A. The cup product $H^2(\Sigma A) \times H^2(\Sigma A) \to H^4(\Sigma A)$ of a suspension is trivial [15, Corollary 4.11] and the naturality of cup product implies that the cup product $H^2(X, A) \times H^2(X, A) \to H^4(X, A)$ is trivial for $j > 0$ as well. This shows that $w \smile w \in B^4(X, A)$. □

B Persistent integral homology computations

B.1 Algorithm for the Earliest Solution problem

The `earliest_solution` algorithm is used to find the persistence of a (co)cycle. This is closely related to computing persistent homology, which is a well-studied problem, at least for coefficient in a finite field. We adapt the boundary matrix reduction algorithm by Edelsbrunner, Letscher and Zomorodian [23], originally developed for persistent homology. Note that this algorithm, unlike classical Gaussian elimination or Smith normal form algorithms, is incremental, which is required by our application. Moreover, it only uses column operations, making efficient implementation for sparse matrices relatively easy.

Efficiency over finite fields. Recent work on computing persistent homology over finite fields resulted in significant performance improvements, see [8, 40]. Despite the cubic worst-case bound, linear scaling is achieved on practical datasets, involving sparse matrices of size $10^9 \times 10^9$ and more. This

encouraged us to adapt the modern version of the classical persistence algorithm to solve our problem over the integers, rather than adapting classical Smith or Hermite normal form algorithms to the persistent setting.

Reduced form and reduction. We adapt the notation common in computational topology literature: the *lowest nonzero* of a nonzero column is defined as the lowest position (largest index) with nonzero coefficient. A sub-matrix is called *reduced* if all the lowest nonzeros are unique. In other words, given a lowest nonzero, there may be other nonzero entries in the same row, but they must not be *lowest* nonzeros. When this invariant is not satisfied, we say there is a collision. By *lowest value*, we refer to the value of the lowest nonzero entry.

The algorithm starts from an empty matrix and adds one column at a time, maintaining the reduced prefix of the matrix, R. The rightmost column of each prefix is called the current column.

Procedure `reduce_column` reduces the column `curr` with respect to the reduced prefix `R`.

```
def reduce_column(current, R, force_divisibility=False):
        while curr collides with some column in R:
                coll = column in R colliding with curr
                P = lowest_value(coll)
                Q = lowest_value(curr)
                if force_divisibility is True and P does not divide Q:
                        return curr
                use Euclid to find nonzero a,b,c,d s.t.
                    gcd(a,b) == 1 and gcd(c,d) == 1 and
                    a*P + b*Q == gcd(P,Q) and
                    c*P + d*Q == 0
                coll, curr = a * coll + b * curr, c * coll + d * curr
        return curr
```

Procedure `earliest_solution` solves the stated problem.

```
def earliest_solution(M, a):
        augment M with identity matrix (above)
        augment a with zeros (above)
        R = []
        for col in M:
                reduced = reduce_column(col, R, force_divisibility=False)
                append reduced to R
                reduce_column(a, R, force_divisibility=True)
                if a is zero:
                        return change of basis of a
        return None
```

Correctness. The basic operation is the addition of two different columns,

without affecting the column span of the relevant matrix prefix. Therefore the solution is unaffected.

To retrieve the solution, at each step we attempt to reduce the input column vector `a` with respect to the currently reduced prefix `R`. The solution exists iff `a` becomes zero, and is encoded by (the negation of) the change of basis column of `a`. Since we solve the equation $Ax = a$ (and not $Ax = ka$) we perform additional divisibility check: see `force_divisibility` in the `reduce_column` procedure.

From finite fields to integers. In the integral case, we may modify both the current column and the colliding column. This is in contrast with the finite field case, in which only the current column is modified. To determine the required linear combination of columns, we use the extended Euclid algorithm. As a result, the lowest *value* of a certain column might change (decrease) many times during the reduction of other columns, but the *position* is fixed once the column is reduced. Because of this, while reducing column `a`, we need to take into account previously reduced columns, and not only the current one. Moreover, after a colliding column is affected, it may not be in the column span of the prefix of the original matrix ending at this column. At this stage, however, this column may only affect columns that succeed the currently reduced column. Therefore, correctnes of the algorithm is unaffected.

Efficiency. For efficiency reasons we use one technique suggested in [8]: The current column is stored in a data-structure handling fast column additions and maximum element queries. One natural choice is a balanced binary search tree; more efficient alternatives are available. This way we avoid the following common bad case: Let n be the total number of columns in M. When an current column becomes dense (the number nonzero entries is $\Theta(n)$), adding $\Theta(n)$ sparse columns takes time quadratic in n. Avoiding this situation does *not* imply that we can efficiently handle matrices that become dense due to fill-in. However in practice, often a small number of columns display this behavior.

We also perform the computations in an on-line fashion – we read columns one by one and stop once a solution is found. This gives significant practical improvements, because often the necessary matrix prefix is very small compared to the matrix of the entire complex.

Overall the algorithm performs well, exhibiting roughly linear scaling in the number of nonzero entries of the original matrix. In particular, we didn't observe coefficient blowup. Note that the lowest nonzero position in the current column decreases after resolving each collision, so the number

of collision resolutions is quadratic in n. Each such resolution requires combining two columns, potentially of size $O(n)$, due to fill-in. This leads to cubic worst case running time bound, assuming that the magnitude of coefficients can be bounded by a constant. Of course, there exist cases when the blowup does occur, and the above analysis does not apply. More advanced algorithms can be used to alleviate the effect of the blowup, possibly at the cost of simplicity and efficiency in the situations that we encountered thus far.

B.2 Algorithm for the Persistent Generators problem

After possibly refining the sequence $(A_i)_i$ we may assume that for each filtration value i there is exactly one $(n-1)$-simplex Δ_i^{n-1} of that filtration value, that is $\Delta_i^{n-1} \in A_{i-1} \setminus A_i$ (and possibly several $(n-2)$-simplices in $A_{i-1} \setminus A_i$). This makes the algorithm and its analysis simpler.

The algorithm for Persistent Generators problem follows. By the statement "reduce the column" we refer to the procedure `reduce_column` from Section B.1

- Let M be the coboundary matrix for $\delta^{n-1}\colon C^{n-1}(X;\mathbb{Z}) \to C^n(X;\mathbb{Z})$, that is, it consists of columns $\boldsymbol{m}$ for integral combination for $\delta\Delta^{n-1}$ for each $\Delta^{n-1} \in X$ of growing filtration values. In exactly the same way we create the coboundary matrix N for δ^{n-2}. During the reduction of the matrix M we keep track of the change of basis. (On low level, we augment the matrix M and then the change of basis vector for a given column is encoded in the augmented part of that column.)

- We initialize $G := ()$ to be an empty sequence of $(n-1)$-cocycles on X and $\mu(-1) := 0$.

- For each filtration value $i = 0, \ldots, h$ do the following:
  - Reduce all columns of N of the filtration value i.
  - Reduce the column $\boldsymbol{m}$ of M of filtration value i (i.e., corresponding to the cochain $\delta\Delta_i^{n-1}$). Let $\boldsymbol{g}$ be the change of basis vector after the reduction.
  - If the reduced column $\boldsymbol{m}$ equals to 0 and there is no reduced column $\boldsymbol{n}$ in N such that $\mathrm{lowest}(\boldsymbol{n}) = i$ and $\boldsymbol{n}_{\mathrm{lowest}(\boldsymbol{n})} | \boldsymbol{g}_{\mathrm{lowest}(\boldsymbol{g})}$ then add $\boldsymbol{g}$ to G and set $\mu(i) := \mu(i-1) + 1$.
  - Otherwise set $\mu(i) := \mu(i-1)$

- Output G and $\mu(1), \ldots, \mu(h)$.

Theorem B.1. *The above algorithm solves the* Persistent Cycles *problem. Namely, after its ith iteration, the cohomology group $H^{n-1}(X, A_i)$ is generated by $[g_1], \ldots, [g_{\mu(i)}]$ where cycles $g_1, \ldots, g_{\mu(i)} \in Z^{n-1}(X, A_i; \mathbb{Z})$ correspond to the columns $\boldsymbol{g}$ added to G during the iterations $0, \ldots, i$.*

Proof. We proceed by induction on the value i. For $i = -1$ the claim trivially holds.

Let us assume that $i \geq 0$ and that the claim holds for $i - 1$. First observe that once the new column $\boldsymbol{m}$ is not reduced to a zero column, then $Z^{n-1}(X, A_i; \mathbb{Z}) = Z^{n-1}(X, A_{i-1}; \mathbb{Z})$. Otherwise $Z^{n-1}(X, A_i; \mathbb{Z}) = Z^{n-1}(X, A_{i-1}; \mathbb{Z}) \oplus \langle g \rangle$ where the cocycle $g \in Z^{n-1}(X, A_i; \mathbb{Z})$ corresponds to the change of basis vector $\boldsymbol{g}$ in the ith iteration. We only have to check whether $[g]$ is a linear combination of $[g_1], \ldots, [g_{\mu(i-1)}]$. By the induction hypothesis, it happens if and only if g is cohomologous to a cocycle in $Z^{n-1}(X, A_{i-1}; \mathbb{Z})$. And this in turn is true if and only if we can reduce to zero the lowest nonzero component (that is, the ith component) of the vector $\boldsymbol{g}$ by adding a combination of columns from N of filtration value at most i (these columns generate the group of coboundaries $B^{n-1}(X, A_i; \mathbb{Z})$). Since this is exactly the reduced part of the matrix N, it is enough to find if there is a column with the lowest nonzero index equal to i and check the divisibility condition. □

C Some details of our implementation: exploiting the cubical structure.

The domain X we consider in the implementation is a cube $[0, 1]^m$ triangulated as follows. We define $I_1, \ldots, I_m$ to be unit intervals subdivided into n_i equidistant intervals of length $1/n_i$, $i = 1, \ldots, m$. This yields a cubical set structure $\prod_j I_j$ on the unit cube. Further, we subdivide each m-cube into $m!$ simplices via the Freudenthal triangulation [5, p. 154]. The resulting triangulation naturally corresponds to the product $\prod I_j$, understood as a product $\prod_j I_j$ of *simplicial sets* [30]. We call the set of vertices a *grid*. A function $f : X \to \mathbb{R}^n$ is then given by a set of $\mathbb{R}^n$-vectors in each vertex (a multidimensional rasterized image), together with a simplexwise Lipschitz constant α.

Many operations in the algorithm outlined in the paper body—such as the computation of the pullback $y = f^\sharp(z)$, computing codifferentials of cochains and the $\smile_{n-3}$ operation—can be done *locally*, without having to

work with the full lists of simplices of a given dimension. The only step where global structure is needed is the EARLIEST SOLUTION algorithm, where the matrix M has rows indexed by the set of all k-simplices and columns indexed by $(k-1)$-simplices. Indeed, the construction of this matrix and computation with it are the most memory- and time-consuming operations, as the number of k-simplices increases rapidly with dimension.

Therefore, we implemented a modification of the algorithm described above based on the more feasible *cubical* structure. Instead of working with the simplicial filtration A_r, we use the filtration $A_r^\square \subseteq A_s^\square$ where $A_r^\square$ is defined to be the triangulation of the set of all cubes c such that $|f(v)| \geq v$ for all vertices v of c. These are still simplicial complexes and f', y, the extension $\tilde{y}$ of y and $\delta\tilde{y}$ are defined with no changes. For the computation of persistence, however, we switch to the cubical setting via the Eilenberg-Zilber reduction [24]. We denote by $C^*(\Pi_j I_j)$ the simplicial cochains and by $\otimes_j C^*(I_j)$ the cubical cochains: there exist chain homomorphisms of degree 0

$$AW : \otimes_j C^*(I_j) \to C^*(\Pi_j I_j)$$
$$EML : C^*(\Pi_j I_j) \to \otimes_j C^*(I_j)$$

such that $EML \circ AW$ is the identity and $AW \circ EML$ is chain homotopic to the identity. Both maps induce cohomology isomorphisms and can be relatively easily implemented using common formulas [35]. This allows us to switch between the simplicial and cubical cochains anytime we need. Within the computation of persistence of the primary obstruction, we compute the smallest j so that there exists a cubical cochain $c_\square \in Z^{n-1}(X, A_{r_j}^\square)$ such that $\delta c_\square = (\delta\tilde{y})_\square := EML(\delta\tilde{y})$. The matrix computation part EARLIEST SOLUTION deals with the cubical coboundary matrix, which is significantly smaller than the simplicial one. For an illustration, the number of k simplices in the triangulation of one m-cube is 50 already for $(m, k) = (4, 2)$ and more than 4 millions for $(m, k) = (8, 6)$.

For the secondary obstruction, we need to convert $c_\square$ back into a simplicial cochain and construct a simplicial $(n-1)$-cochain $x \in \Omega(A_{r_j})$ that extends y on A_{r_j} and is a global cocycle. We denote by SHI the cochain homotopy map $SHI : C^*(\Pi_j I_j) \to C^{*-1}(\Pi_j I_j)$, satisfying $AW \circ EML - \mathrm{id} = \delta \circ SHI + SHI \circ \delta$. Then, using $\delta^2 = 0$, we obtain

$$\delta AW c_\square = AW \delta c_\square = AW(\delta\tilde{y})_\square = AW\, EML\, \delta\tilde{y} =$$
$$= \delta AW\ EML\, \tilde{y} = \delta\tilde{y} + \delta\, SHI\, \delta\tilde{y}.$$

We compute $\tilde{c} := AW(c_\square) - SHI(\delta\tilde{y})$ which is a *simplicial* cochain zero on $A^\square_{r_j}$ and the last equation asserts that $\delta\tilde{c} = \delta\tilde{y}$ which allows us to compute the simplicial cocycle

$$x := \tilde{y} - \tilde{c} \in \Omega(r_j)$$

having avoided to work with large lists of all simplices. The computation of $v(x) := (x \mod 2) \smile_{n-3} (x \mod 2)$ is done on the simplicial level and the property $x \in Z^{n+1}(X, A^\square_{r_j}; \mathbb{Z}_2)$ depends on the fact that the chosen ordering of vertices of X and Σ^{n-1} is compatible with f'. Then we again apply the EML operator to $v(x)$ and convert it to a cubical cochain $v(x)_\square$. Persistent generators $w \in Z^{n-1}(X, A^\square_r; \mathbb{Z})$ and their Steenrod images can be computed on the cubical level and the cubical persistence of the secondary obstruction is done via a cubical coboundary matrix, this time in dimension $n, (n+1)$.

There is a price to pay for the more convenient cubical filtration: it is courser than the simplicial one and we can no more use the estimate on the robustness of zero set derived in Section 3. On one hand, we have the relation $A^\square_r \subseteq A_r$. This implies that whenever $f'|_{A_r}$ is homotopic to $f/|f|$, then so is $f'|_{A^\square_r}$, and non-extendability of $f'|_{A^\square_r}$ to higher skeleta of X implies the non-extendability of $f'|_{A_r}$. Thus whenever the algorithm certifies non-extendability of $f|_{A^\square_r}$ and $r > \alpha n^{1/p}$, then $r - \alpha$ is a lower bound on the robustness of zero. If we can additionally prove that $A_s \subseteq A^\square_r$, then the extendability of $f'|_{A^\square_r}$ implies the extendability of $f'|_{A_s}$. The simplexwise α-Lipschitz condition implies that any two vertices u, v of a cube c satisfy $|f(u) - f(v)| \leq 2\alpha$, as u, v can be connected by path of at most two simplicial edges. This implies $A_{r+2\alpha} \subseteq A^\square_r$ and extendability of $f'|_{A^\square_r}$ to all of X implies the upper bound $r + 3\alpha$ on the robustness.

Based on the experience with various functions in low dimensions, the practical performance has not one significant bottleneck. The most resources-consuming steps in the primary obstruction computation include the computation of the cubical filtration and the EARLIER SOLUTION subroutine, especially if the matrix has millions of columns.

D Experiments with formulas.

Our prototype is implementation in Python using numpy. While the efficiency is severely limited by this choice, we were able to run examples for $\dim X \leq 8$.

Although the strength of our methods is primarily in cases where we have uncertainty on f, the following examples provide some observations

about the performance in cases where the function values at the vertices are computed using exact formulas.

In what follows, we assume a subdivision of the interval $[-1,1]$ into a set I_g of g equidistant points and consider a grid $I_g^m \subseteq [-1,1]^m$ of size g^m. We chose the max-norm ℓ_∞ on $\mathbb{R}^n$ which yields the smallest value of $\alpha n^{1/p}$ and allows smaller "initial r_0". This is desired, as for small grids it often happens that α is large and $\alpha n^{1/p}$ may be larger then the robustness of zero in which case we fail to detect anything.

Testing primary obstruction on a quadratic function. First we used the function $f : [-1,1]^n \to \mathbb{R}^n$ given by

$$\begin{aligned} f_1(x) &= x_1^2 - x_2^2 - \ldots - x_n^2 \\ f_2(x) &= 2x_1x_2 \\ &\ldots \\ f_n(x) &= 2x_1x_n. \end{aligned}$$

This function has a single zero in the origin and has been used as a benchmark example in [26]. If n is even, then the zero is robust and its robustness equals $\min_{x\in\partial[-1,1]^m} |f(x)|$: the zero has index two in this case. If n is odd, then the zero has index zero and can be removed by arbitrary small perturbations. Using the max-norm, we use calculus to derive the estimate $|f(x) - f(y)|_\infty \leq 2n\,|x-y|_\infty$. Whenever x, y are contained in one simplex of the triangulation, then the max-norm satisfies $|x-y|_\infty \leq 2/(g-1)$ and we can use the piecewise Lipschitz constant $\alpha := 4n/(g-1)$.

For our experiments, we now assume that f is only given via a list of function values in the grid g^n and the simplexwise Lipschitz constant $\alpha = 4n/(g-1)$. In this case, we have $\dim X = n$ so only primary obstruction can be nontrivial. Using Lemma 2.4 for $p = \infty$ yields that whenever $r > \alpha$, then f' is simplicial on A_r and homotopic to $f/|f|$. The relation $A_r^\square \subseteq A_r$ then implies that f' is simplicial and homotopic to $f/|f|$ on $A_r^\square$ as well. Thus we can run our algorithm with an initial $r_0 = \alpha = 4n/(g-1)$, compute the persistence of primary obstruction $r_1 \geq r_0$ and—using the estimates at the end of Section C—conclude that whenever $r_1 > r_0$ then the robustness of zero is between $r_1 - \alpha$ and $r_1 + 3\alpha$.

The following table illustrates some properties of the computations: n refers to the dimension, g is the number of points subdividing the intervals in each dimension, the initial $r_0 = \alpha$ is as above, r_1 is the persistence of the primary obstruction, δ^{n-1} the number of columns of the cubical coboundary matrix $C_\square^{n-1} \to C_\square^n$, the "computed robustness" displays lower and upper

bounds $r_1 - \alpha$, $r_1 + 3\alpha$ on the robustness of zero and the "true robustness" column the approximation of the real robustness of zero of the function f defined exactly by the above formula,[12] all wrt. the $|\cdot|_\infty$ norm.

n	g	$r_0 := \alpha$	r_1	# columns of δ^{n-1}	computed robustness	true robustness
2	20	0.421	0.8643	760	$[0.443, 2.127]$	0.828
2	100	0.08	0.8285	19 800	$[0.748, 1.07]$	0.828
2	500	0.016	0.83	499 000	$[0.814, 0.878]$	0.828
3	20	0.632	$= r_0$	21 660	≤ 2.379	0
3	50	0.245	$= r_0$	360 150	≤ 0.971	0
3	100	0.121	$= r_0$	2 940 300	≤ 0.483	0
4	20	0.842	$= r_0$	548 720	≤ 3.37	0.667
4	30	0.552	0.711	2 926 680	$[0.159, 2.367]$	0.667
4	40	0.41	0.667	9 491 040	$[0.257, 1.897]$	0.667

The total running times of these 9 computations is displayed in Figure 3. The data suggest that the at least in these simple cases, the running time is approximately linear in the number of $(n-1)$-simplices. The $n = 3$ case is below the other two, because the primary obstruction is trivial there and the EARLIEST SOLUTION subroutine terminates almost immediately, using only very few columns in the matrix reduction.

In higher dimensions, the condition $r_0 \geq \alpha := 4n(g-1)$ is too strict and we can only hardly start with r_0 that is smaller than the robustness. However, we could take r_0 to be the minimal value for which $f'|_{A^\square_{r_0}}$ is simplicial: in all cases that we have tried, we verified that f was nowhere zero on $A^\square_{r_0}$ with f' homotopic to $f/|f|$ on $A^\square_{r_0}$.

Surprisingly, starting with minimal r_0 for which $f'|_{A^\square_{r_0}}$ is simplicial and computing the minimal r_1 for which f' is extendable, yields in many cases a much better estimate of the robustness of zero set than the estimates based on Lemma 2.4: this can already be seen in the above table where the r_1 column is a good approximation of the true robustness. We give the table with smaller grids that continues to higher dimensions.

[12] This is 0 in odd dimensions and $\frac{2}{\sqrt{n+1}}$ in even dimension.

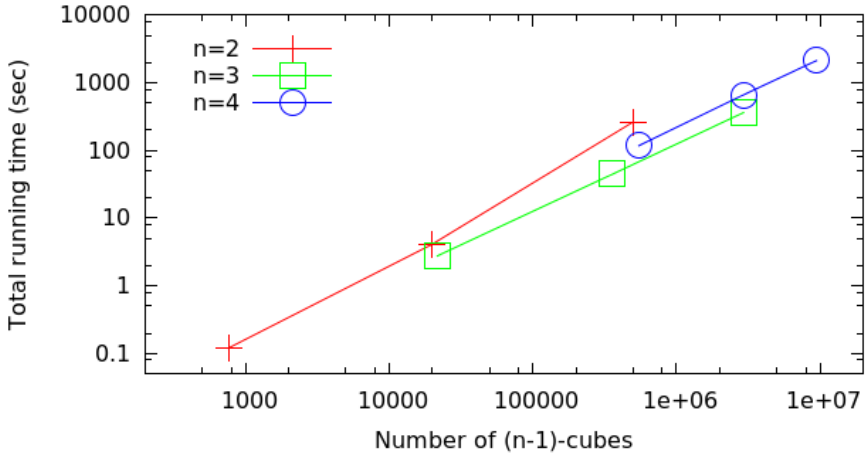

Figure 3: Total running time of the computations, performed on Intel Xeon CPU X5680, 3.3 GHz, as a function of the number of $(n-1)$-simplices (which equals the number of columns in the coboundary matrix).

n	g	min simplicial r_0	r_1	true robustness
2	10	0.025	0.889	0.828
3	10	0.025	$= r_0$	0
4	10	0.025	0.667	0.667
5	10	0.074	$= r_0$	0
6	10	0.074	0.667	0.58
7	6	0.241	$= r_0$	0
8	5	0.251	1.0	0.522

It is an interesting question to find natural conditions on functions, other than the relatively weak simplexwise Lipschitz property, that justify the usage of smaller grids and guarantee that the robustness of zero set is close to the computed minimal r for which $f'|_{A_r^\square}$ becomes extendable.

A function with nontrivial secondary obstruction. A second function we experimented with is $h : [-1,1]^{n+1} \to \mathbb{R}^n$ given by

$$\begin{aligned} h_1(x) &= 2x_0x_2 + 2x_1x_3 \\ h_2(x) &= 2x_1x_2 - 2x_0x_3 \\ h_3(x) &= x_0^2 + x_1^2 - x_2^2 - x_3^2 \\ h_4(x) &= x_4 \\ &\dots \\ h_n(x) &= x_n. \end{aligned}$$

The restriction of h to $\partial[-1,1]^{n+1}$ is the generator of the nontrivial homotopy group $\pi_n(S^{n-1})$: in case $n = 3$ it is the Hopf map and for $n > 3$ its iterated suspension. The robustness of zero equals $\min_{x\in\partial[-1,1]^{n+1}} |h(x)|$ and[13] it is the simplest example where a nontrivial secondary obstruction occurs. Common tests for zero verification such as the degree test would fail here.

Again, we work with the max-norm ℓ_∞ and derive, via elementary calculus, the estimate $2(n+1)$ on the global Lipschitz constant. This yields $|h(x)-h(y)| \le \frac{4(n+1)}{g-1} =: \alpha$. Assume that only α and a list of function values in a grid g^{n+1} is given. If g is large enough so that $r_0 := \alpha$ is smaller than $\min_\partial |h(x)|_\infty = \sqrt{3}-1$, then the algorithm computes the minimal $r_1 > r_0$ for which $h|_{A_r^\square}$ is extendable to a nowhere zero function $[-1,1]^{n+1} \to \mathbb{R}^n \setminus \{0\}$. For this to succeed, we need to take g at least 22 for $n = 3$.

Again, the program gives surprisingly good results for much smaller grids, although Lemma 2.4 gives no guarantees. In all cases that we tried,

[13]The robustness is 1 in the ℓ_2 norm and $\sqrt{3}-1$ in the max-norm.

we verified that whenever r_0 is large enough so that $h'|_{A_r^{\square}}$ is simplicial, then it is homotopic to $h/|h|$ and the algorithm computes that the secondary obstruction dies close to the real robustness of zero.

n	g	min simplicial r_0	r_1	true robustness
3	10	0.1	0.79	0.732
4	10	0.111	0.79	0.732
5	8	0.163	0.816	0.732